



⑯ Europäisches Patentamt
European Patent Office
Office européen des brevets



⑪ Publication number:

0 525 986 A2

⑫

EUROPEAN PATENT APPLICATION

⑯ Application number: 92305993.5

⑮ Int. Cl. 5: G09G 1/16, G09G 5/36

⑯ Date of filing: 29.06.92

⑯ Priority: 26.07.91 US 736686

⑯ Applicant: SUN MICROSYSTEMS, INC.

⑯ Date of publication of application:
03.02.93 Bulletin 93/05

2550 Garcia Avenue
Mountain View, CA 94043(US)

⑯ Designated Contracting States:
DE FR GB IT NL

⑯ Inventor: Priem, Curtis

4052 Kettering Terrace
Fremont, California 94536(US)

Inventor: McIntyre, Bruce

20090 Rodriguez Avenue, Nr. 8E
Cupertino, California 95014(US)

Inventor: Malachowsky, Chris

3775 ThrushWay

Santa Clara, California 95051(US)

Inventor: Moffat, Guy

181A El Carmelo

Palo Alto, California 94306(US)

⑯ Representative: Wombwell, Francis
Potts, Kerr & Co. 15, Hamilton Square
Birkenhead Merseyside L41 6BR(GB)

⑯ Apparatus for fast copying between frame buffers in a double buffered output display system.

⑯ An output display system including a first frame buffer; a second frame buffer; apparatus for transferring data from the second frame buffer to an output display device; apparatus for controlling the transfer of data to each of the frame buffers, this last-men-

tioned apparatus including apparatus for writing new data only to the first frame buffer, and apparatus for both reading data from the first frame buffer and writing the data to the second frame buffer during the same operation.




FIG. 2

EP 0 525 986 A2

BACKGROUND OF THE INVENTION

Field Of The Invention

This invention relates to computer output display apparatus and, more particularly, to methods and apparatus for eliminating frame tearing from a computer output display through the use of inexpensive double buffering.

History Of The Prior Art

A typical computer system generates data which is displayed on an output display. This output display is typically a cathode ray tube which produces a number of full screen images one after another so rapidly that to the eye of the viewer the screen will appear to display constant motion when a program being displayed is capable of producing such motion. In order to produce the individual images (frames) which are displayed one after another, data is written into a frame buffer. The frame buffer stores information about each position on the display which can be illuminated (each pixel) to produce the full screen image. For example, a display may be capable of displaying pixels in approximately one thousand horizontal rows each having approximately one thousand pixels. All of this information in each frame is written to the frame buffer before it is scanned to the display.

When data describing an entire picture exists in the frame buffer, the frame may be transferred to the display. Typically, data is transferred from the frame buffer to the display pixel by pixel and line by line beginning at the upper left hand corner of the display and proceeding horizontally from left to right line by line downwardly to the lower right hand corner of the display. In order for the picture to appear continuous on the output display, the successive frames in the frame buffer must be constantly scanned to the output display at a rate of thirty frames per second or more.

While each frame of data is being scanned to the display, new data to appear in the next frame must be transferred to the frame buffer. In general, only data which is changing replaces old data in the frame buffer at a position representing that pixel position on the screen. All unchanged data remains in the frame buffer without change. New data to be displayed in a frame may be written to any portion of the frame buffer at any time. In order to allow information to be both written to the frame buffer and scanned out of the frame buffer to the output display simultaneously, two ported video random access memory (VRAM) is used for the frame buffer. Data is written through one port and scanned to the display through the other. VRAM is more expensive than typical dynamic access ram

(DRAM) because providing the two ports requires a significant number more transistors.

If data is being placed in a VRAM frame buffer at the same time that information is being scanned to the display, it is possible that information being scanned to the display will come from two succeeding time displaced frames. For example, if scanning is proceeding at a faster rate than data is being written to the frame buffer and a portion of the frame buffer which is changing (being written) is scanned to the display, a portion of the display will be from what should be a first frame and a portion from what should be a second later frame. The display of portions of two time displaced frames simultaneously is called frame tearing. This can be disconcerting where the information is rapidly changing as in real time video, for images may be grossly distorted.

In order to eliminate frame tearing, double buffering is used. Double buffering uses two complete frame buffers each of which may store one entire frame. Data is written to one frame buffer and scanned to the display from the other. In its simplest form, this is accomplished using a pair of VRAM frame buffers and multiplexing the data in one or the other to the display. In this form, data is never written to the frame buffer during the time its data is being scanned to the display. Once a frame has been completely written, it may in turn be scanned to the display and all further data written to the other frame buffer. Since data is never written to a frame buffer while its contents are being scanned to the display, frame tearing cannot occur.

This simple form of double buffering is somewhat expensive because it uses two entire VRAM frame buffers and includes control signal generating circuitry and a multiplexor for switching between the two frame buffers.

One of the primary aims of computer designers is to allow a number of individual programs to run on a computer and be displayed simultaneously on an output display of that computer. Typically, when a number of individual programs are displayed on a computer output display, each individual program appears in a window, typically a rectangular area of the screen which may be moved about, enlarged and reduced in size, and otherwise manipulated. If a number of programs can be run and displayed in a number of windows simultaneously, the work being accomplished using the computer may be accelerated.

Typically, information being written to the individual windows by the different individual programs will be written at different rates. For example, the information being directed to a window displaying real time video changes very rapidly while the information typed from the keyboard to a

word processor program being displayed in another window changes much more slowly. Consequently, the rate at which frames are changed varies from program to program.

The simplest form of double buffering described above is very useful when a single program is being run on the output display. However, where a number of programs are being run simultaneously in different windows on the same output display, this form of double buffering is insufficient. The reason for this is that the simple form of double buffering requires that the entire contents of each frame buffer be scanned to the display. If data is being written to a number of windows at asynchronous rates, then the timing at which writing occurs differs from window to window; and it is very difficult to adjust the timing of the writing so that writing does not occur to a frame buffer being scanned to the display. To solve this problem, an advanced form of double buffering has been used which adds another buffer called a window identification (ID) plane. The window identification plane provides a storage position for each pixel displayed on the output display. Stored in these positions of the window ID plane are identifications of the window to which each pixel of data is related. Use of this plane allows pixels from any buffer to be selected for display at any time. Thus, the window ID plane may be used to scan to the display data from any window to which data is not being written at the time of the scan. Thus, this form of double buffering allows frame tearing to be eliminated where multiple active windows appear simultaneously on the output display.

This second form of double buffering is quite expensive because it not only uses two entire VRAM frame buffers and circuitry for controlling and multiplexing from the two frame buffers to the display, it also adds an ID plane containing memory for each pixel of the display and circuitry for selecting pixels to be displayed based on the windows in which they appear.

Experimenters have looked for arrangements which would reduce the expense of the two VRAM buffers and the control circuitry for multiplexing between the two VRAM frame buffers or between individual pixels in the first and second forms of double buffering. One form of double buffering which has been used in the prior art to reduce cost replaces one of the VRAM frame buffers with a single-ported DRAM frame buffer and eliminates the control circuitry for scanning to the display from either of the frame buffers. Instead, all frames are scanned to the display from the single VRAM frame buffer; and all new data is written to the DRAM frame buffer. Once written, the data in the DRAM frame buffer is transferred by the central processing unit from the DRAM frame buffer to the

5 VRAM frame buffer. This requires a read of the DRAM by the processor followed by a write of the data to the VRAM frame buffer. Typically this takes place in blocks of data as large as the processor is able to transfer on its bus (e.g., thirty-two bits); and the transfer is repeated over and over until complete.

10 This form of double buffering is much less expensive than the other forms because a less expensive DRAM replaces one of the VRAM frame buffers and the control circuitry for multiplexing is eliminated. This arrangement is also useful because it works well with software conforming to the X11 standard (X Windows) which does not expect to see more than a single frame buffer and stores information to be transferred to the frame buffer in a section of main memory. To this software, the DRAM frame buffer appears to be a portion of main memory. The arrangement also offers the advantage that individual windows may be transferred from the invisible DRAM frame buffer to the VRAM frame buffer since the central processing unit may selectively control the areas to be transferred.

15 20 25 30 35 However, the transfer of information from the DRAM frame buffer to the VRAM frame buffer is relatively slow compared to the typical scan rate to the display. Consequently, it is possible for writing to the VRAM frame buffer to occur in a position from which information is being scanned to the display; and the problem of frame tearing arises. Were the writing to proceed at a rate faster than the scan out to the display, it would be possible to eliminate this problem by writing below the row being scanned and still utilize inexpensive circuitry to accomplish double buffering.

Summary Of The Invention

40 It is, therefore, an object of the present invention to reduce the expense of double buffering arrangements of the prior art while retaining the ability to eliminate frame tearing.

45 It is another more specific object of the present invention to increase the speed of transfer between frame buffers in a system in which information from only one of two frame buffers is scanned to the output display.

50 55 These and other objects of the present invention are realized in an output display system comprising a first frame buffer; a second frame buffer; means for transferring data from the second frame buffer to an output display device; means for controlling the transfer of data to each of the frame buffers, this last-mentioned means including means for writing new data only to the first frame buffer, and means for both reading data from the first frame buffer and writing the data read to the second frame buffer during the same operation.

These and other objects and features of the invention will be better understood by reference to the detailed description which follows taken together with the drawings in which like elements are referred to by like designations throughout the several views.

Brief Description Of The Drawings

Figure 1 is a block diagram illustrating a prior art arrangement for reducing the cost of double buffering arrangements.

Figure 2 is a block diagram of an arrangement in accordance with the present invention for providing double buffering.

Notation And Nomenclature

Some portions of the detailed descriptions which follow are presented in terms of symbolic representations of operations on data bits within a computer memory. These descriptions and representations are the means used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. The operations are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like. It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities.

Further, the manipulations performed are often referred to in terms, such as adding or comparing, which are commonly associated with mental operations performed by a human operator. No such capability of a human operator is necessary or desirable in most cases in any of the operations described herein which form part of the present invention; the operations are machine operations. In all cases the distinction between the method operations in operating a computer and the method of computation itself should be borne in mind. The present invention relates to apparatus and to method steps for operating a computer in processing electrical or other (e.g. mechanical, chemical) physical signals to generate other desired physical signals.

Detailed Description Of The Invention

Referring now to Figure 1, there is illustrated a circuit 10 constructed in accordance with the prior art. The circuit 10 includes only the rudiments of the circuitry required to provide data to an output display terminal used in a typical computer system. Other portions necessary for providing the operations of a computer are well known to those skilled in the art and are not shown in the figure. Illustrated are a central processing unit 12 which 5 may control the operation of the entire computer system and which represents in Figure 1 circuitry for providing data to be display on an output display 14. In order to accomplish the transfer of the data from the central processing unit 12 to the output display 14, first and second frame buffers 16 and 17 are utilized.

In the simplest form of double buffering described above, data is written from the central processing unit 12 to one frame buffer and scanned to the display 14 from the other. In its simplest form, this is accomplished using a pair of VRAM frame buffers and multiplexing the entire frame of data in one of the frame buffers 16 or 17 to the display by means of a multiplexor 19. The 10 data transferred by the multiplexor 19 is converted from digital to analog form by a digital-to-analog converter 20 and scanned to the display. In this form of double buffering, data is never written to a frame buffer 16 or 17 during the time data is being scanned to the display from that frame buffer. 15 Once a frame has been completely written to a frame buffer 16 or 17, the data in that frame buffer may in turn be scanned to the display; and new data may be written to the other frame buffer. 20 Since data is never written to a frame buffer while its contents are being scanned to the display, frame tearing cannot occur.

The implementation of this form of double buffering is somewhat expensive because it uses two frame buffers both of which are constructed of 25 expensive VRAM. Moreover, the arrangement requires control signal generating circuitry for selecting the frame buffer the information from which is to be multiplexed for display and a multiplexor 19 for switching between the two frame buffers 16 and 17.

In order to reduce these expenses, one prior 30 art arrangement has replaced the VRAM used in the frame buffer 16 with DRAM. Since DRAM is single ported, the frame buffer 16 does not provide an output which may be scanned directly to the display 14. Consequently, the lines from the frame buffer 16 to the multiplexor 19 (which are dotted in the figure) are eliminated. Since no output is transferred from the frame buffer 16 to the multiplexor 19, the multiplexor 19 (again shown in dotted outline) is also eliminated. With no multiplexor, the 35 control circuitry for selecting one or the other of the

frame buffers 16 or 17 to scan to the display 14 is also unnecessary and is eliminated. This substantially reduces the cost of the system.

All new data is written by the central processing unit 12 to the frame buffer 16. All data is scanned to the display 14 from the frame buffer 17. Once a frame stored in frame buffer 16 has been changed, the frame is transferred to the frame buffer 17. To accomplish this transfer, a control circuit 23 within the central processing unit 12 selects a portion of the frame buffer 16 to be read, typically thirty-two bits or some amount usually equivalent to the width of the bus. This data is read and latched into the central processing unit 12. This read typically requires three clock periods for row addressing and four clock periods for column addressing by the control circuit 23. The central processing unit 12 then writes the information read from the frame buffer 16 to the frame buffer 17. This again typically requires three clock cycles for row addressing by the control circuit 23 in frame buffer 17 and three clock cycles for column addressing in that frame buffer. The additional clock period for column addressing in reading frame buffer 16 during the read operation is actually necessary in order to provide a dead cycle on the bus between the reading by the central processing unit 12 of the frame buffer 16 and the following write cycle to the frame buffer 17 so that two devices are not accessing the bus simultaneously.

Thus, a total of thirteen or more clock cycles are required to transfer a given amount of information from the frame buffer 16 to the frame buffer 17 in this arrangement. This transfer is repeated a sufficient number of times to transfer the desired amount of data (which may be as large as an entire frame) from the frame buffer 16 to the frame buffer 17. As may be visualized, this is a relatively slow process and, when copying entire frames, allows only approximately twenty frames to be transferred per second. A typical display, on the other hand, may be receiving information from the frame buffer 17 at a rate of seventy-six frames per second. The scan proceeds at approximately three times the rate of the write to the frame buffer 17.

Consequently, this less expensive arrangement allows the scan to the display to catch up with the writing of data into the frame buffer 17 from the frame buffer 16 so that frame tearing may occur.

To obviate the problem of frame tearing in this less expensive double buffering system, the arrangement of the present invention illustrated in Figure 2 has been devised. The arrangement 25 includes a central processing unit 27, a first frame buffer 28 which may be constructed of VRAM, a second frame buffer 29 which may be constructed of DRAM, a digital-to-analog converter 31, and an output display 33. The arrangement 25 functions in

the same general manner as does the least expensive arrangement of Figure 1. That is, all new data is written by the central processing unit 27 to the DRAM frame buffer 29. All data is scanned to the display 33 from the VRAM frame buffer 28. Once new data has been written to a frame stored in frame buffer 29, the frame is transferred to the frame buffer 28. To accomplish this transfer, the central processing unit 27 reads a selected portion of the frame buffer 29, typically thirty-two bits, and writes that data to the frame buffer 28. This process is accomplished over and over until the desired amount of data has been transferred.

In contrast to the time taken to transfer data from the frame buffer 16 to the frame buffer 17 in the arrangement of Figure 1, the arrangement of the present invention transfers the data approximately four times as fast. Thus, approximately eighty frames per second may be written to the frame buffer 28 from the frame buffer 29; and tearing of frames scanned to the display may be eliminated.

The fast copying of data from the frame buffer 29 to the frame buffer 28 is accomplished in the following manner. Rather than the single control circuit 23 for controlling the selection of data to be accessed in both of the frame buffers 16 and 17 as in the arrangement of Figure 1, the arrangement 25 includes within the central processing unit 27 (or other device controlling the rendering into and reading from the frame buffers 28 and 29) a pair of individual control circuits 34 and 35. The first of these circuits 34 controls the accessing of the frame buffer 29, and the second circuit 35 controls the accessing of the frame buffer 28. By utilizing individual control circuits 34 and 35, both frame buffers 28 and 29 may be accessed simultaneously. Data is still written only to the frame buffer 29 by the central processing unit when rendering. However, when data is to be read from the frame buffer 29 and written to the frame buffer 28, the control circuit 34 selects the appropriate row and column addresses in the frame buffer 29, and the control circuit 35 selects the same row and column addresses in the frame buffer 28. Then the control circuit 34 reads the accessed data in the frame buffer 29 and places it on the bus where the information is written to the same accessed addresses in the frame buffer 28. The data is not latched into the central processing unit 27; and, consequently, no dead cycle is needed for bus turn around so that two devices are not attempting to access that bus simultaneously.

This arrangement saves a drastic amount of access time. First, there are not independent read and write cycles so that row and column selects for the two frame buffers occur simultaneously and occupy half the time. Then, the clock cycle for bus

turn around is not needed. Finally, since the control circuitry need not address first one frame buffer and then another, it may essentially latch the row address at the first transfer. This eliminates the time for succeeding row accesses for the remainder of any particular row to be displayed. Consequently, after the first piece of data is accessed, the remaining transfers from a row each require only a single three clock access time to both read the frame buffer 29 and write the data to the frame buffer 28. This is less than one quarter of the time required by the prior art circuitry.

Because the time necessary to copy from the frame buffer 29 to the frame buffer 28 is reduced to less than one quarter, more than four times as many full frames of data can be written to the frame buffer 28 using the invention as in prior art double buffering arrangements. Thus, over eighty frames per second may be written to the frame buffer 28 while the highest scan rate from the frame buffer is seventy-six frames per second. At this rate, the scan cannot catch up with the writing to frame buffer 28 and frame tearing will not occur. In order to assure that copying does not catch up with the scan, however, it is necessary to begin all copying between the frame buffer 29 and 28 at a point to be displayed below the point at which scanning is taking place. This is simply accomplished by referring to the row position of the scan in the circuitry controlling the scan of information from the frame buffer 28 to the display.

Although the present invention has been described in terms of a preferred embodiment, it will be appreciated that various modifications and alterations might be made by those skilled in the art without departing from the spirit and scope of the invention. The invention should therefore be measured in terms of the claims which follow.

Claims

40

1. An output display system comprising a first frame buffer; a second frame buffer; means for transferring data from the second frame buffer to an output display device; means for controlling the transfer of data to each of the frame buffers, this last-mentioned means including means for writing new data only to the first frame buffer, and means for both reading data from the first frame buffer and writing the data to the second frame buffer during the same operation.

45

2. An output display system as claimed in Claim 1 in which the means for both reading data from the first frame buffer and writing the data read to the second frame buffer during the same operation comprises control means for

50

55

simultaneously selecting addresses for row accesses of the first and second frame buffers, and for simultaneously selecting addresses for column accesses of the first and second frame buffers.

5

10

15

20

25

30

35

40

45

50

55

3. A method of controlling the transfer of data to an output display system comprising the steps of transferring new data to be displayed to a first frame buffer; reading data to update the display from the first frame buffer and simultaneously writing that data to a second frame buffer; and transferring a frame of data from the second frame buffer to an output display.

4. A method of controlling the transfer of data to an output display system as claimed in Claim 3 in which the step of reading data to update the display from the first frame buffer and simultaneously writing that data to a second frame buffer further comprising the steps of simultaneously selecting addresses for row accesses of the first and second frame buffers, and simultaneously selecting addresses for column accesses of the first and second frame buffers.




FIG. 1




FIG. 2